

**IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF MICHIGAN
SOUTHERN DIVISION**

RYAN MILLIRON,

Plaintiff,

v.

UNITED STATES DEPARTMENT OF
DEFENSE,

Defendant.

Civil Action No. 1:23-cv-1222

Hon. Robert J. Jonker
U.S. District Judge

Hon. Phillip J. Green
U.S. Magistrate Judge


EXHIBIT 2
Final Response to DOD FOIA Request
23-F-1597

Fancy Bear/APT28 Attribution Analysis

Summary Data

ASTROLAVOS LAB

Georgia Institute of Technology



APT28 Attribution Analysis: Summary Data

August 7, 2016

CONTENTS

1 Introduction	1
2 Internet (b)(4) Analysis	2
2.1 (b)(4) Analysis	2
2.2 Analysis Based on Rhamnousia Framework	5
3 (b)(4) Analysis	12
3.1 (b)(4)	15
3.2 (b)(4) Analysis	17
4 Analysis and Recommendations	17
4.1 Additional Comments	18
5 Conclusion	18

Abstract

We provide detailed attribution data analysis of the July 2016 Fancy Bear/APT28 campaign. Our unique contributions include: analysis of the (b)(4), observations about the (b)(4) used by attacks, and indicators of the (b)(4) by (b)(4) including (b)(4).

INTRODUCTION

Although attack attribution is a developing field, there are several proven early techniques [10, 17, 23]. The Georgia Institute of Technology (“Georgia Tech”) has worked on attribution problems for a decade, and has the benefit of countless failures and a few successes. The Astrolavos Lab¹ at Georgia Tech is starting a large-scale, long-term attribution research effort (under DARPA’s Enhanced Attribution project), to combine machine learning techniques, large data analysis, and innovative feature selection. Called the *Rhamnousia Framework*, this project combines (b)(4) with next-generation analysis.

Using some of our early attribution techniques, this report examines the Fancy Bear/APT28 campaign. This campaign is apparently related to the recent cyber attack on the Democratic National Committee (“DNC”) and is widely attributed to Russian national efforts [2, 19].

Using the (b)(4) information as a starting point, we (b)(4) significantly. Further, we were able to point to (b)(4) identities and (b)(4) the APT28 (b)(4) (b)(4). Further, our report identifies additional (b)(4) which may be related to APT28, but otherwise appear malicious.

¹<http://astrolavos.gatech.edu/>




Figure 1: Information flow used in our analysis. Using the (b)(4) we (b)(4) (b)(4)

The specific steps in our analysis are summarized in Figure 1. We first perform a (b)(4) analysis, by noting (b)(4). We then perform an attribution analysis using (b)(4); (b)(7)(E) (b)(4). Using this expanded set, we then perform a timing analysis, to identify other (b)(4). On this (b)(4), we then repeat the attribution analysis, to identify any (b)(4). While each step may add a small amount of false positives, our approach hopefully eliminates any possible false negatives, yielding the largest possible set of indicators.

The unique contributions of our report include:

- A more complete identification of indicators of compromise.
- Analysis of the (b)(4) to create malicious (b)(4) used for APT28 and other attacks. This analysis also includes (b)(4).
- A (b)(4) analysis of (b)(4) which yields a smaller set of (b)(4) for future analysis.

Our conclusion also offers some (perhaps obvious) suggestions to those who have more resources for a deeper investigation. (E.g., we note various companies which likely have relevant records, which would require some proper paperwork or authority to pursue.)

INTERNET (b)(4) ANALYSIS

(b)(4) ANALYSIS

We start with (b)(4) [8, 20, 22] which list three (b)(4) used as indicators for the “Fancy Bear/APT28” attacks. These three (b)(4) appear in Figure 2. For the rest of this document we will call these (b)(4) as:

- the (b)(4) identified by (b)(4) [8].
- the (b)(4) identified by (b)(4) [20].
- the (b)(4) identified by (b)(4) [22].

The (b)(4) in these three (b)(4) according to these (b)(4) were used in the Fancy Bear/APT28 campaign. Using these (b)(4) we attempt an attribution analysis on (b)(4) (b)(4) and (b)(4). To start, we first validate these intelligence sources.

Since the (b)(4) in the (b)(4) are largely overlapping, we decided to extensively analyze the most compact (b)(4). The very first question we will have to answer is: Do these (b)(4) exhibit (b)(4) that suggest targeted (and potentially nation state) attacks? For this, we first used (b)(4) from one component of the Rhamnousia Framework, namely (b)(4) and focused on (b)(4) behind the (b)(4). We looked for (b)(4) (b)(4) starting with January 2015 (b)(4).



Figure 2: The three sets of (b)(4) obtained from (b)(4) and the fourth (b)(4); (b)(7)(E); (b)(3); 50 U.S.C. § that we were able to discover through the Rhamnusia Framework.

If these (b)(4) were used solely in a targeted attacks, then there should be little interest in (b)(4).




Figure 3: The frequency distribution (b)(4)




Figure 4: The cumulative distribution function (CDF) of (b)(4).



Figure 5: The geographic distribution of (b)(4)




Figure 6: The geographic distribution of (b)(4)




Figure 7: (b)(4) volumes since 2015.




Figure 8: (b)(4) volume in (b)(4)

Figures 3 and 4 show the results of the (b)(4) analysis per (b)(4). It is clear that not all the (b)(4) were used in the attack. The very simple reason behind this is that some of



Figure 9: The three [redacted] that the Rhamnousia Attribution process separated [redacted]. Using the [redacted] as a seed, the Rhamnousia framework was able to derive the [redacted].

them clearly have been [redacted] (b)(4). Such [redacted] (b)(4) do not follow the signature of a targeted attack.

ANALYSIS BASED ON RHAMNOUSIA FRAMEWORK

Historically, the authors of this report have studied how large and small-scale infections look like from the point of DNS [redacted] (b)(4) [4]. Using statistical properties from such past research efforts, [redacted] from the Rhamnousia Framework (such as [redacted] from [redacted] [redacted] and [redacted] from [redacted] we tried to [redacted] (b)(4) which could explain their public classification as APT indicators. The result of this [redacted] (b)(4) process can be seen in Figure 9. These three different groups exhibit significantly different resolution behavior. In particular, the [redacted] (b)(4) in [redacted] (b)(4) are very popular, and are resolved in networks all over the world. These, we suspect, are really unlikely to be true indicators of an APT attack.

The next [redacted] (b)(4) could very well be part of an APT campaign. This is because they are being looked up from [redacted] (b)(4) in central Europe and Middle East. Given the discussion around the APT28 attack from Fireeye², we suspect that these set of domain names could reflect activities around this operation.

Finally, we are left with an extremely interesting [redacted] (b)(4) namely [redacted] (b)(4) in Figure 9. This [redacted] (b)(4) have little in common with the [redacted] (b)(4) in [redacted] (b)(4) in terms of [redacted] (b)(4). These [redacted] (b)(4) are [redacted] (b)(4) from [redacted] (b)(4) in US, Europe and [redacted] (b)(4) Brazil. [redacted] (b)(4) can be clearly seen in Figure 6. On the other hand, the rest of the [redacted] (b)(4) result in [redacted] (b)(4) patterns reflected in Figure 5. It is clear, just from the graphic, that the [redacted] (b)(4) is very likely to be part of an APT campaign — likely different to the rest of the campaigns in the [redacted] (b)(4).

The [redacted] (b)(4) are most notable, and we decided to dive deeper into its [redacted] (b)(4). In Table 1 we can see the [redacted] (b)(4) worldwide, from which we saw [redacted] (b)(4) from the beginning

²<https://www.fireeye.com/blog/threat-research/2014/10/apt28-a-window-into-russias-cyber-espionage-operations.html>

(b)(4)

(b)(4)

(b)(4); (b)(6)

Table 1: (b)(4)

(b)(4); (b)(6)

Table 2: (b)(4)

of January 2015 until the end of July 2016. At this point we should remind the reader that (b)(4) (b)(4) been publicly attributed to the Fancy Bear/APT28 campaign by (b)(4) [8]. If this is true, we should expect that we will see (b)(4) close to the public disclosure of the “Fancy Bear” attack, in June of 2016. We tried to validate this by looking at their historic (b)(4) especially will the respect of the most interesting (b)(4) in Table 1; the (b)(4) (b)(4)

The (b)(4) have (b)(4) data for the (b)(4) listed in Table 2. These appear to be the (b)(4) and/or (b)(4) for any illicit activity behind (b)(4) (b)(4). The (b)(4) in the table appear malicious, and our preliminary analysis let us to suspect that very likely are Russian connections (b)(4) in Germany and Turkey. Often, former FSB and Russian Business Network operations used (b)(4)³ We should also note that the (b)(4) (b)(4) which likely means that it has been used as the primary infection vector for the attack. Finally, the (b)(4) related to (b)(4) and (b)(4)

we believe it is likely a (b)(4) (b)(4). After this analysis, we decided to investigate the (b)(4) traffic and the (b)(4) in (b)(4) even further. At this point we have enough indicators to believe that they could have been used attacks against

³<http://www.computerworld.com/article/2532289/cybercrime-hacking/cyberattacks-knock-out-georgia-s-internet-presence.html>

the (b)(4) network, and likely (b)(4) networks too. Using the methodology described in Figure 1 behind our attribution framework, we were able to (b)(4). These (b)(4) (b)(4) can be seen in Figure 2 and we will call (b)(4). In the same figure you can see a comparative analysis between the (b)(4) and the (b)(4) (b)(4). In all case, the (b)(4) contains between 45 and 47 (b)(4) names that **are currently unknown to the rest of the community**. We highly recommend their use as indicators of compromise, and recommend operators use (b)(4) to check for initial signs of infections.

The (b)(4) to a variety of different servers around the globe. Figure 12, shows the directed graph between the (b)(4) and (b)(4) (with their country flag) that they point into. As we can see from Table 4, we have seven (b)(4) in (b)(4) six in (b)(4) five in (b)(4) three in (b)(4) two in (b)(4) and (b)(4) and finally one in (b)(4) Romania, France and Bulgaria.



Figure 10: The temporal patterns behind (b)(4) from (b)(4) in the (b)(4) network.

very first major spike (b)(4) appears to happen in the (b)(4) in the early days of February 2016. The spike in (b)(4) from the (b)(4) network matches the very first significant spike for all (b)(4). The only explanation we can give on this given the data we have access to is that (b)(4) appears to be among the first (if not the first) places that were targeted by the "Fancy Bear" campaign.

Following this observation, we immediately went back in our Framework and we looked for (b)(4) (b)(4) occurring withing a few seconds time window, with the respect of the (b)(4) of (b)(4) in the (b)(4). In Figure 10, we can see two different cases of (b)(4) behind the (b)(4) in the (b)(4). In the CASE 1, we see how they orchestrate (b)(4). In the CASE 2, we case see (highlighted with with red) (b)(4) immediately after the resolutions behind the (b)(4). While we cannot obviously be sure that these resolutions were omitted by the

The last experiment we conducted is a (b)(4) analysis of the (b)(4) in (b)(4). In Table 7, we can see the (b)(4) patterns observed by the Rhamnousia Framework, for the entire (b)(4). We see a growth pattern that makes a lot of sense. In 2015, the overall (b)(4) volumes are low, but as the various campaigns take action, we see the growth of (b)(4). This growth peaks in June/July of 2016, where the indicators of compromise became public from the various threat reports (i.e., [8,20,22]). Now, if we isolate the (b)(4) in (b)(4) (see Figure 2), we observe something completely different. First, the bulk of (b)(4) for these (b)(4) (b)(4) took place in February 2016 and not in June/July. This can be explained because most of the (b)(4) in (b)(4) are not known to the community. Thus, when the public reports came online, these (b)(4) were (b)(4) (b)(4).

However, the most important observation from Figures 7 and 8, is that the

same (potentially infected) host, it is quite possible that if an infection was successfully delivered in the (b)(4) network, the (b)(4) and (b)(4) to be used as a pivot point to (b)(4) networks. This makes the forensic investigation of historic DNS resolutions for all domains in the (b)(4) very important.




Figure 11: (b)(4) analysis between (b)(4) and (b)(4)

If (b)(4) can be identified, we should examine their (b)(4), (b)(7)(E), (b)(3):50 U.S.C. § 3024(i)(1) patterns. If they have (b)(4) and/or (b)(4) that resembles (b)(4) (b)(4) those could be considered as suspicious.

After analyzing the (b)(4) from (b)(4) we were able to identify (b)(4) that fit this pattern. The most interesting from them was (b)(4), which (b)(4)

This appears to be (b)(4); (b)(6) Oddly, (b)(4) (b)(4) and not (b)(4)

Our first reaction was to try and find whether (b)(4) It is common practice to have (b)(4) for a variety of reasons (content delivery, load balancing, localization, etc.). Looking at the (b)(4), we can see that this is a (b)(4) Looking a bit closer into the temporal patterns of this (b)(4) see Figure 11, we can see that the DNS resolutions for this (b)(4) are both constant and periodic (note green circular dot in the figure 11).

This (b)(4) activity is clearly very interesting, simply because we cannot see any benign activity (given the (b)(4), that would justify this periodic (likely automaton-driven) (b)(4) (b)(4). Thus, our first reaction was to find more about all (b)(4) (b)(4) To analyze the behavior of this zone we conducted a week's long experiment, for the first week of March. During this week we observed the following:

- The query types we saw in this week followed the following (b)(4) pattern: (b)(4) (b)(4) Clearly an unusual pattern for (b)(4) that used for (b)(4) The large volume of (b)(4) denote system processes (i.e., (b)(4) trying (b)(4) This can occur, for example, when a (b)(4) without performing (b)(4) (b)(4) The (b)(4) (usually (b)(4) makes note of (b)(4) but avoids (b)(4)



(b)(4);
(b)(7)(E);
(b)(3); 50
U.S.C. §
3024(i)(1)

Figure 12: The (b)(4) behind (b)(4) that support (b)(4)

listing the (b)(4) instead (b)(4). This phenomena is observable even at (b)(4) (b)(4) [6], associated with (b)(4)

- For the A record types, we saw 47 different (b)(4) with the (b)(4) (b)(4) (b)(4) (b)(4) are generally not popular according to their (b)(4) (b)(4)

Most interestingly, an overall (b)(4) from 109 (b)(4) looked up these (b)(4). The first five (b)(4) were: (b)(4). Looking a bit closer into the (b)(4) we saw that two of them were: (b)(4) (b)(4) and (b)(4). Our suspicion is these (b)(4) reflect activities from within (b)(4). Table 3 lists the exact (b)(4) that these (b)(4) were (b)(4). The (b)(4) behavior cannot be explained by any normal user interaction. The two very sensitive (b)(4) networks, and the remaining residential address space makes this *at best* very interesting (b)(4) lookup pattern. What kind of software or Internet activity could result in such a (b)(4)—one that interests only a few (b)(4) around the world, plus (b)(4) and (b)(4).

(b)(4)

(b)(4); (b)(6)

(b)(4)

Table 3: [redacted] in the US that [redacted] (b)(4) for the [redacted] (b)(4)

(b)(4); (b)(6); (b)(7)(E); (b)(3):50 U.S.C. § 3024(i)(1)

Table 4: (b)(4) the (b)(4) in the

(b)(4);
(b)(7)(E);
(b)(3):50
U.S.C. §
3024(i)(1)

(b)(4);
(b)(7)(E);
(b)(3)-50
U.S.C. §
3024(i)(1)**(b)(4) ANALYSIS**

Since many of the (b)(4) shared a (b)(4) via (b)(4) (see Figure 14), we decided to look closely at the (b)(4) that created (b)(4).

Previously, cyber criminals created ICANN accredited registrars to easily generate new command-and-control domains, and obfuscate domain ownership. Following the 2008 revocation of Estdomains [24], and 2013 shutdown of Dynamic Dolphin [16], ICANN began auditing accredited registrars and verifying the accuracy of domain WHOIS data [12].

Despite criticism of this enforcement effort [1], the ICANN scrutiny made operation of an explicitly criminal-run or criminal-friendly registrar far more risky. As a result, many criminal groups switched to domain reseller services—registration services that are not directly ICANN accredited, but who license their Top Level Domain (TLD) access through an affiliate model.

While this still provides cyber criminals with DNS agility, it does not afford them complete WHOIS privacy. Granted, “privacy protected” domain registrations are offered by criminal-friendly resellers. But the non-obfuscated/non-private customer data is maintained (as per ICANN escrow regulations [11]) in the registrar’s Extensible Provisioning Protocol (EPP) access system⁵. Thus, while privacy protected domain information is common in malicious C2 domains, the accredited registrar has the customers’ true credentials: non-public contact email, login IPs, default NS settings, and often credit card information.

Thus, even if the domain reseller is criminal-friendly or criminally operated, the otherwise private WHOIS information for C2 domains is held by third parties, who has real incentives to (a) help remediate cyber crime when notified; and (b) have significant incentives to never facilitate cyber crime through their accredited company. For the most part, this means domain registrars will blacklist or ban abusive domains in their reseller programs. In other cases, it means the registrar will revoke the privacy rights of the customer (often without notice, depending on the AUP).

Our Framework includes (b)(4) allowing us insight into (b)(4) below (b)(4). We used this to identify the owners of the APT28 (b)(4) and also obtained the (b)(4) used for (b)(4) and (b)(4). No doubt the (b)(4) and the (b)(4) are (b)(4) but this reveals additional information about the APT28 (b)(4). We know, for example, that under the ICANN regulations the (b)(4) must have worked (and for some (b)(4) even the (b)(4) must work).

Figure 13 shows the (b)(4) organized by (b)(4), which also appear in Appendix A. The (b)(4) for the (b)(4) are also noted, and colorized (yellow/dashed (b)(4)). (b)(4) were merely used for (b)(4). Note also that the (b)(4) are used only once. In fact, only one (b)(4) is used for (b)(4)—an indication of significant operational security not usually seen in mere criminal groups.

We next looked at the (b)(4) to determine if it was merely a (b)(4). Previously, (b)(4) used (b)(4) but this exposed (b)(4) to additional (b)(4). (b)(4) now resells only (b)(4) and processes payments through (b)(4), a (b)(4).

We note a few facts about (b)(4) which illustrate its role in facilitating APT attacks:

- Approximately 13% of the (b)(4) entire portfolio in (b)(4) – 1,012 out of (b)(4) total (b)(4) are suspended for reported abuse. A very high ratio relative to other (b)(4).
- A complete list of the portfolio of (b)(4) owned by (b)(4) appears as Appendix 5. We note that this includes several (b)(4) with (b)(4). Overall, the portfolio is very small for a (b)(4) making their operation likely unprofitable.

⁵<http://tools.ietf.org/html/rfc5731>

⁶<https://bitcointalk.org/index.php?topic=711690.0>

⁷<https://en.wikipedia.org/wiki/BitPay>

(b)(4);
(b)(7)(E);
(b)(3)-50
U.S.C. §
3024(i)(1)(b)(4);
(b)(7)(E);
(b)(3)-50
U.S.C. §
3024(i)(1)

(b)(4)

(b)(4);
(b)(7)(E);
(b)(3)-50
U.S.C. §
3024(i)(1)



Figure 13: Overview of APT28, (b)(4) and (b)(4) Yellow/dashed (b)(4) (all (b)(4) indicate (b)(4)

- We also note that (b)(4) purchased their own (b)(4) through their (b)(4) (b)(4) [13]. This is a risky (b)(4) since any dispute (policy, billing, etc.) with the (b)(4) could jeopardize the (b)(4) site itself. When combined with the very high abuse rate, this appears to be a risk for the (b)(4) operation, assuming it wishes to continue as (b)(4)
- Based on a few (b)(4) (proxied through (b)(4), the (b)(4) for (b)(4) is BitPay, an Atlanta-based company. They accept no other form of payment.
- In terms of ICANN compliance, (b)(4) appear to permit (b)(4) absent even customer verification. They do ultimately send verification emails to customers, but these are not barriers to registration. As case could be made for non-compliance and revocation of reseller status.
- For a registrar with (b)(4) (ever) and only a handful (b)(4) their site nonetheless has language offerings in Russian and English. This evidently reflects the needs of their target customers, and the skill sets of the operators.
- Many hundreds of domains created through (b)(4) are phishing related, and security companies have noted (b)(4) [21].

(b)(4) Against all this, we decided to try and find the operators of (b)(4) When the (b)(4) site is administered, the owners login (b)(4) in Sri Lanka, mostly from (b)(4) Table 5 shows a few details about these logins. We urge further inquiry into these (b)(4) which may be (without other factual basis) to be proxies. In Appendix 5, we list all the (b)(4) used by the owners in the last month. We note that one (b)(4) from Wisconsin (b)(4); (b)(6) is unusual in this set. It is classified as a cable customer by the (b)(4) [9], and was not listed on any behavioral-based (b)(4) (e.g., Tor exit nodes, proxy lists, spam/infected hosts lists, etc.) It is not known if the Wisconsin host is a proxy, or a negligent/non-proxied error during site administration.

The (b)(4) operators have a youtube channel [7], with a single video ⁸, which has just over 9,000 views and two channel subscribers. A still from the video appears as Figure 15. An reverse image search for these people turns up a “testimonial site” ⁹ which produces paid video endorsements of products (usually fraudulent, SEO-related, or high yield investment scams). Figure 16 shows an image of the testimonial site, and also indicates the actors in the (b)(4) site. Note that of all the actors on the site, the (b)(4) actors are the only ones identified by ‘code names’ (e.g., “Orange and OG”) rather than normal user names (which nonetheless might be fake, but are normal names such as “Dave”). Archived

(b)(4)




Figure 14: (b)(4)



Table 5: (b)(4) for administration of (b)(4)

Figure 15: Putative owners of (b)(4) (from [7])

copies of the site show (b)(4) (the (b)(4) first appeared as early as Feb. 19, 2012 on the site ¹⁰, and the couple first appeared in archives dated March 20, 2012 ¹¹. Since (b)(4) was registered in February of 2012, (and the (b)(4) video published Jun 22, 2012), the (b)(4) may have some social or business connection to the (b)(4) domain owners. I.e., it is our theory that the earliest (b)(4) might have some relation to the (b)(4) owners. We urge cautious further investigation into (b)(4) to explore these possibilities.

The (b)(4) operators also have an idle twitter account, with a few dozen followers. They all appear to be merely (b)(4) enthusiasts, but further inquiry of their twitter followers is required.


(b)(4) HOSTING

We note that (b)(4) makes use of (b)(4) and (b)(4). During the relevant period (late 2015 to present), (b)(4) published (b)(4).



The site also published (b)(4) for (b)(4) a business operation outsourcing company in India. We suggest: perhaps the US-based companies (b)(4) Google) may have records of the site verification and web hosting sign up, which may yield (b)(4) and (b)(4).





1 of 2

07/31/2016 06:06 AM

Figure 16: (b)(4) from (b)(4) promotional video.

Table 6: (b)(4) prices during creation of (b)(4)

(b)(4) ANALYSIS

As noted, the (b)(4) accepts only (b)(4). Lacking a (b)(4) for the (b)(4) of the APT28 (b)(4) (e.g., (b)(4)), we could instead examine the (b)(4) (b)(4) for clues, around the time key APT28 (b)(4). For example, the (b)(4) (b)(4) was first seen in the (b)(4) on 2015-12-11 10:43:58 -0000.

Lacking a precise price for (b)(4) registrations at that time period, one might survey history (b)(4) prices during that period. Appendix 5 shows the dot com prices on that date for a few (b)(4) (b)(4). We recommend using only the low end of the scale (b)(4) since the (b)(4) (b)(4) competes in the (b)(4) using (b)(4).

Between the hours of 0900 and 1045, the price for (b)(4) on Dec 11, 2015 is listed in Table 6.

Thus, one could examine all (b)(4) from approximately 10:15:00 to 10:45:00, December 11, 2015, in amounts corresponding to approximately \$5 to \$10. The payment processor for (b)(4) is BitPay, and it was likely used on December 11, 2015 (but this cannot be demonstrated). We are not aware of an existing (b)(4) database indexed by date, but will endeavor to create such a data store during our project. In followups to this report, we will look into the finances of (b)(4). But others who may have a fully indexed (b)(4) can use this analysis to perhaps identify the (b)(4) used in the APT (b)(4).

ANALYSIS AND RECOMMENDATIONS

Here, we provide cautious, factual analysis, and recommendations for further inquiry.

- We note that the youtube channel for (b)(4) has two subscribers. Given the operational security demonstrated by the (b)(4) we would be surprised if the subscribers are linked to (b)(4). This may be a dead-end (e.g., the subscribers could be (b)(4)) but proper paperwork served on Youtube and Gmail would provide answers. At best, it may yield another (b)(4) and (b)(4) through Google.
- We note that (b)(4) and (b)(4) would provide more detail about the (b)(4) interactions with the (b)(4). We urge that logging be implemented, if such logs are not presently collected.
- We recommend that all (b)(4) created by the (b)(4) be candidates for blacklisting, and that an NS-level RPZ be implemented, where possible. Appendix B lists all the (b)(4) we can discover using authorities associated with (b)(4) or (b)(4).
- The owners of (b)(4) logged in from (b)(4) Wisconsin, but usually used (b)(4) in Sri Lanka. This should be looked into. It may be an accidental failure to use Sri Lankan proxies, or it may be breach of their systems.

- We note that (b)(4) uses BitPay, a US-based company (3405 Piedmont Rd NE, Atlanta, GA 30305 (855) 424-8729), for payment clearance. This company is presumably responsive to paperwork, and may have contact information for (b)(4) that are not yet discovered.
- The (b)(4) was previously used by (b)(4). They are part of (b)(4) a (b)(4) company in (b)(4) Utah (b)(4) (b)(4). They might have contact information for the (b)(4) operators, and presumably would respond to paperwork.
- The (b)(4) site was hired by the (b)(4) to produce a video. The (b)(4) company has an archived (internet archive, Feb. 2011) phone number of (b)(4) (b)(4) area code from 2012), and a live/current phone number of (b)(4) (b)(4) area code). The (b)(4) company often promotes SEO, high-yield investment scams, and other (likely fraudulent) campaigns. Given their connection to abuse, we are skeptical they have records of the 2011 (b)(4) video production, or would be a wise target for normal service of paperwork. (Since they either participate in fraud, or knowingly make money off fraudulent campaigns, they would seem just as likely to share with others any paperwork they receive.) We suggest cautious inquiry.

ADDITIONAL COMMENTS

While focused narrowly on APT28 (b)(4) this study ultimately identified attributes of an attack on an electoral system. Our study provided more information about the attack origins, but we stop short of noting direct connections to specific operators. Our future research project, starting soon under DARPA, will provide such connections. But for now our (b)(4) analysis provides (we believe) unique insights, which may help others tasked with attribution analysis.

Assuming there are connections to Russia, we note the long history of Russian interference with elections. It therefore might be useful to consider the following scenarios:

- One should expect continued use of sockpuppets social accounts (“web brigades”), to influence public perceptions in the US elections [14, 15].
- One might anticipate the insertion of fake emails into dumps of legitimate emails, thereby leaving the afflicted parties unable to effectively deny the validity or origin of the messages.
- One should anticipate targeted DDoS attacks against political websites, and infrastructure essential for voting, to cast doubt on the results. This has been used in previous Russian-directed cyber attacks to affect the outcome of elections [18].
- One should consider that the visible DNC and DCCC attacks behind the Fancy Bear/APT28 campaigns were perhaps not the first nor the only targets — given the network behavior observed at (b)(4) and (b)(4) (just by our network visibility). We advice proper authorities to use the indicators of compromise in the (b)(4) for post mortem forensic analysis using historic passive DNS traces from their organizations.

Again, we note that this portion of the report is beyond our scope of expertise. But the (b)(4) analysis we have provided suggests enormous resources are being used to penetrate sensitive networks, and we urge others to consider future possible attacks.

CONCLUSION

Our previous work in attribution focused on pooling intelligence and coordinating investigations, on an ad hoc basis. We have identified the key (b)(4) necessary to build a science of attribution

(b)(7)(E),
(b)(3) 50
U.S.C. §
3024(i)(1)

(b)(4)
(b)(7)(E),
(b)(3) 50
U.S.C. §
3024(i)(1)

(b)(7)(E),
(b)(3) 50
U.S.C. §
3024(i)(1)

(b)(4),
(b)(7)(E),
(b)(3)-50
U.S.C. §
3024(i)(1)

(b)(4),
(b)(7)(E),
(b)(3)-50
U.S.C. §
3024(i)(1)

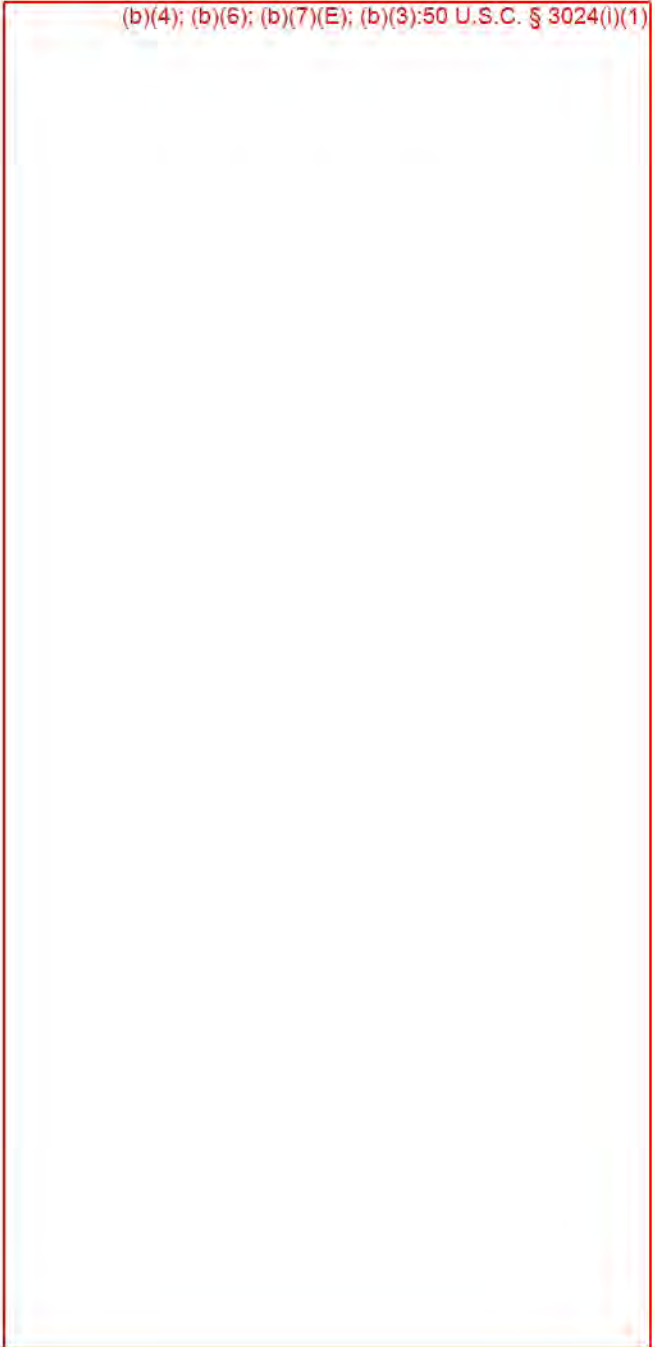
and traceback. Our plans for the Rhamnousia Framework are [REDACTED]
[REDACTED] and innovate new [REDACTED] (b)(4) methods to attribute
network abuse. We hope that our preliminary analysis of the Fancy Bear/APT28 campaign yields additional
insights.

APPENDIX: (b)(4) AND (b)(4) RECORDS FOR (b)(4) APT28
(b)(4)

(b)(4); (b)(6); (b)(7)(E); (b)(3); 50 U.S.C. § 3024(i)(1)

(b)(4); (b)(6); (b)(7)(E); (b)(3); 50 U.S.C. § 3024(i)(1)

(b)(4); (b)(6); (b)(7)(E); (b)(3):50 U.S.C. § 3024(i)(1)




APPENDIX: (b)(4) FOR (b)(4)

(b)(4); (b)(6); (b)(7)(E); (b)(3); 50 U.S.C. § 3024(i)(1)

(b)(4); (b)(6); (b)(7)(E); (b)(3); 50 U.S.C. § 3024(i)(1)

(b)(4); (b)(6); (b)(7)(E); (b)(3); 50 U.S.C. § 3024(i)(1)



(b)(4); (b)(6); (b)(7)(E); (b)(3); 50 U.S.C. § 3024(i)(1)



(b)(4); (b)(6); (b)(7)(E); (b)(3):50 U.S.C. § 3024(i)(1)

APPENDIX: [REDACTED] (b)(4) FOR [REDACTED] (b)(4)

List of [REDACTED] (b)(4) registered domains (including those not yet in TLD zone).

(b)(4); (b)(7)(E); (b)(3):50 U.S.C. § 3024(i)(1)

(b)(4); (b)(7)(E); (b)(3):50 U.S.C. § 3024(i)(1)



(b)(4); (b)(7)(E); (b)(3); 50 U.S.C. § 3024(i)(1)




(b)(4); (b)(7)(E); (b)(3); 50 U.S.C. § 3024(i)(1)




(b)(4); (b)(7)(E); (b)(3); 50 U.S.C. § 3024(i)(1)

(b)(4); (b)(7)(E); (b)(3); 50 U.S.C. § 3024(i)(1)



(b)(4); (b)(7)(E); (b)(3); 50 U.S.C. § 3024(i)(1)

(b)(4); (b)(7)(E); (b)(3); 50 U.S.C. § 3024(i)(1)



(b)(4); (b)(7)(E); (b)(3); 50 U.S.C. § 3024(i)(1)



APPENDIX: POSSIBLE (b)(4) FOR (b)(4)

(b)(4); (b)(7)(E); (b)(3); 50 U.S.C. § 3024(i)(1)

REFERENCES

- (b)(4)
- [2] Dmitri Alperovitch. Bears in the midst: Intrusion into the democratic national committee. <https://www.crowdstrike.com/blog/bears-midst-intrusion-democratic-national-committee/>, June 2016.
- [3] Manos Antonakakis, Roberto Perdisci, David Dagon, Wenke Lee, and Nick Feamster. Building a dynamic reputation system for DNS. In *the Proceedings of 19th USENIX Security Symposium (USENIX Security '10)*, 2010.
- [4] Manos Antonakakis, Roberto Perdisci, Wenke Lee, Nikolaos Vasiloglou, and David Dagon. Detecting malware domains in the upper DNS hierarchy. In *the Proceedings of 20th USENIX Security Symposium (USENIX Security '11)*, 2011.
- [5] Manos Antonakakis, Roberto Perdisci, Yacin Nadjji, Nikolaos Vasiloglou, Saeed Abu-Nimeh, Wenke Lee, and David Dagon. From Throw-Away Traffic to Bots: Detecting the Rise of DGA-Based Malware. In *the Proceedings of 21th USENIX Security Symposium (USENIX Security '12)*, 2012.
- (b)(4)
- [10] Jeffrey Hunker, Robert Hutchinson, and Jonathan Margulies. Attribution of cyber attacks on process control systems. In *Critical Infrastructure Protection II*, pages 87–99. Springer, 2009.
- [11] ICANN. Registrar accreditation agreement. <https://www.icann.org/resources/pages/ra-agreement-2009-05-21-en#3.4.1>, May 2009.
- [12] ICANN. Accuracy. <https://whois.icann.org/en/accuracy>, 2013.
- [13] ICANN. Registrar ids. <http://www.iana.org/assignments/registrar-ids/registrar-ids.xhtml>, July 2016.
- [14] John Kelly, Vladimir Barash, Karina Alexanyan, Bruce Etling, Robert Faris, Urs Gasser, and John Palfrey. Mapping russian twitter. https://cyber.law.harvard.edu/sites/cyber.law.harvard.edu/files/Mapping_Russian_Twitter_2012.pdf, 2012.
- [15] Olga Khazan. Russia's online-comment propaganda arm. <http://www.theatlantic.com/international/archive/2013/10/russias-online-comment-propaganda-army/280432/>, 2013.
- [16] Brian Krebs. Spam-friendly registrar 'dynamic dolphin' shuttered. <http://krebsonsecurity.com/2013/11/spam-friendly-registrar-dynamic-dolphin-shuttered/>, November 2013.
- [17] Thomas Rid and Ben Buchanan. Attributing cyber attacks. *Journal of Strategic Studies*, pages 1–34, 2014.

- [18] Hal Roberts and Bruce Etling. Coordinated ddos attack during russian дума elections. <http://blogs.harvard.edu/idblog/2011/12/08/coordinated-ddos-attack-during-russian-duma-elections/>, December 2011.
- [19] David Sanger and Eric Schmitt. Spy agency consensus grows that russia hacked d.n.c. <http://www.nytimes.com/2016/07/27/us/politics/spy-agency-consensus-grows-that-russia-hacked-dnc.html>, July 2016.
- [20] ThreatConnect Research Team. Fancy bear has an (it) itch that they can't scratch. <https://www.threatconnect.com/fancy-bear-it-itch-they-cant-scratch/>, July 2016.
- [21] ThreatConnect. What's in name server. <https://www.threatconnect.com/whats-in-a-name-server/>, July 2016.

(b)(4)

- [23] David A Wheeler and Gregory N Larsen. Techniques for cyber attack attribution. Technical report, DTIC Document, 2003.

(b)(4)